

Investigating the Characteristics of One-Sided Matching Mechanisms Under Various Preferences and Risk Attitudes

Hadi Hosseini · Kate Larson · Robin Cohen

Received: date / Accepted: date

Abstract One-sided matching mechanisms are fundamental for assigning a set of indivisible objects to a set of self-interested agents when monetary transfers are not allowed. Two widely-studied randomized mechanisms in multiagent settings are the Random Serial Dictatorship (RSD) and the Probabilistic Serial Rule (PS). Both mechanisms require only that agents specify ordinal preferences and have a number of desirable economic and computational properties. However, the induced outcomes of the mechanisms are often incomparable and thus there are challenges when it comes to deciding which mechanism to adopt in practice.

In this paper, we first consider the space of general ordinal preferences and provide empirical results on the (in)comparability of RSD and PS. We analyze their respective economic properties under general and lexicographic preferences. We then instantiate utility functions with the goal of gaining insights on the manipulability, efficiency, and envyfreeness of the mechanisms under different risk-attitude models. Our results hold under various preference distribution models, which further confirm the broad use of RSD in most practical applications.

Keywords One-Sided Matching · Random Serial Dictatorship · Probabilistic Serial Rule · Strategyproofness · Social Welfare · Fairness · Risky Attitudes

1 Introduction

One-sided matching mechanisms have been extensively adopted in many resource allocation settings such as assigning dormitory rooms or offices to students, stu-

Hadi Hosseini
Department of Computer Science, Rochester Institute of Technology, Rochester, NY, USA
E-mail: hhvcs@rit.edu

Kate Larson
Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, Canada
E-mail: klarson@uwaterloo.ca

Robin Cohen
Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, Canada
E-mail: rcohen@uwaterloo.ca

dents to public schools, college courses to students, organs and medical resources to patients, and members to subcommittees [5, 16, 34, 43]. Two prominent (randomized) matching mechanisms that only elicit ordinal preferences from agents are *Random Serial Dictatorship* (RSD) [2] and *Probabilistic Serial Rule* (PS) [14]. Both mechanisms have important economic properties and are practical to implement. The RSD mechanism has strong truthful incentives but guarantees neither efficiency nor envyfreeness. PS satisfies efficiency and envyfreeness; however, it is susceptible to manipulation. Therefore, there are subtle points to be considered when deciding which mechanism to use. For example, given a particular preference profile, the mechanisms often produce random assignments which are simply incomparable and thus, without additional knowledge of the underlying utility models of the agents, it is difficult to determine which is the “better” outcome. Furthermore, properties like efficiency, truthfulness, and envyfreeness can depend on whether there is underlying structure in the preferences, and even in general preference models it is valuable to understand under what conditions a mechanism is likely to be efficient, truthful, or envyfree as this can guide designers choices.

In this paper, we study the comparability of PS and RSD when there is only one copy of each object, and analyze the space of all preference profiles for different numbers of agents and objects. Working in the space of general ordinal preferences, we provide empirical results on the (in)comparability of RSD and PS and analyze their respective economic properties. We show that despite the inefficiency of RSD, the fraction of random assignments at which PS stochastically dominates RSD vanishes, especially when the number of agents is less than or equal to the available objects. We also investigate the manipulability of PS and show that PS is almost always 99% manipulable for any combination of agents and objects, and the fraction of strongly manipulable profiles goes to one as the ratio of objects to agents increases. We show similar trends under lexicographic preferences, and further present results on envy of RSD. Our results show that although the fraction of envious agents grows with the number of agents, there is a sudden drop in the fraction of envious agents when there are equal number of agents and objects.

In Section 5, we instantiate utility functions for agents to gain deeper insights on the manipulability, social welfare, and envyfreeness of the two mechanisms under different risk attitudes. Our main result is that under risk aversion, the social welfare of RSD is as good as PS but RSD does create envy among the agents (though the fraction of envious profiles and the total envy are small). Moreover, when the number of agents and objects are equal, RSD assignments are less likely to be dominated by PS, and overall RSD assignments create negligible envy among agents. We also show that PS is highly susceptible to manipulation in almost all combinations of agents and objects. The fraction of manipulable profiles and the gain from manipulation rapidly increases, particularly when agents become more risk averse. In Section 7, we consider two statistical preference distribution models, namely Mallows Models and Polya-Eggenberger Urn Models, and show that the same patterns and trends hold for various combinations of agents and objects, when varying risk parameters and utility functions.

2 Preliminaries

In this section, we describe the basic one-sided matching problem and introduce the two mechanisms we study in detail, Random Serial Dictatorship (RSD) [2] and Probabilistic Serial Rule (PS) [14]. We then introduce a number of properties and criteria used to evaluate these mechanisms.

A one-sided matching problem consists of a set of n agents, N , and a set of m indivisible objects, M .¹ Each agent $i \in N$ has a private strict preference ordering, \succ_i , over M where $a \succ_i b$ indicates that agent i prefers to receive object a over object b . We represent the preference ordering of agent i by the ordered list of objects $\succ_i = a \succ_i b \succ_i c$ or $\succ_i = (abc)$, for short. We let \mathcal{P} denote the set of all complete and strict preference orderings over M . A *preference profile* $\succ \in \mathcal{P}^n$ specifies a preference ordering for each agent, and we use the standard notation $\succ_{-i} = (\succ_1, \dots, \succ_{i-1}, \succ_{i+1}, \dots, \succ_n)$ to denote preferences orderings of all agents except i and thus $\succ = (\succ_i, \succ_{-i})$.

The goal in a one-sided matching problem is to assign the objects in M to the agents in N according to preference profiles, under the constraint that no object can be assigned to more than one agent. If $m = n$ then this means that each agent will receive exactly one object, however if $m < n$ then some agents will receive no object and if $m > n$ then some agents may receive multiple objects. An assignment is represented as a matrix

$$A = \begin{pmatrix} A_1 \\ A_2 \\ \vdots \\ A_n \end{pmatrix} = \begin{pmatrix} A_{1,1} & A_{1,2} & \dots & A_{1,m} \\ A_{2,1} & A_{2,2} & \dots & A_{2,m} \\ \vdots & \vdots & \ddots & \vdots \\ A_{n,1} & A_{n,2} & \dots & A_{n,m} \end{pmatrix}$$

where $A_{i,j} \in [0, 1]$ is the probability that agent i is assigned object j . We let \mathcal{A} denote the set of all *feasible* assignments where an assignment $A \in \mathcal{A}$ is *feasible* if and only if $\forall j \in M, \sum_{i \in N} A_{i,j} = 1$. If $A \in \mathcal{A}$ is such that $A_{i,j} \in \{0, 1\}$ then we say that A is a *deterministic* assignment; otherwise, A is a *random* assignment. Every random assignment can be represented as a convex combination of deterministic assignments [49], and thus we view random assignments as a probability distribution over a set of deterministic assignments.

2.1 Matching Mechanisms

In general, a *matching mechanism*, \mathcal{M} , is a mapping from the set of preference profiles, \mathcal{P}^n to the set of feasible assignments, \mathcal{A} . That is, $\mathcal{M} : \mathcal{P}^n \mapsto \mathcal{A}$. We focus our attention on two widely studied mechanisms for one-side matching: Random Serial Dictatorship (RSD) [2] and Probabilistic Serial Rule (PS) [13].

RSD relies on the concept of priority orderings over agents. Such an ordering is an ordered list of agents where the first agent gets to select its most preferred object from the set of objects, the second agent then selects its most preferred

¹ This problem is sometimes called the assignment problem or house allocation problem in the literature.

object from the set of remaining objects and so on until no objects remain.² Given a preference profile $\succ \in \mathcal{P}^n$, RSD returns an assignment $RSD(\succ) \in \mathcal{A}$ which is a uniform distribution over all deterministic assignments induced from all possible priority orderings over the set of agents. RSD has been widely adopted for fair and strategyproof assignments for the school choice problem, course assignment, house allocation, and room assignment [1–3, 47]

PS treats objects as a set of divisible goods of equal size and simulates a simultaneous eating algorithm. Each agent starts “eating” its most preferred object, all at the same rate. Once an object is gone (eaten away) then the agent starts eating its next preferred object among the remaining objects. This process terminates when all objects have been “eaten”. Given a preference profile $\succ \in \mathcal{P}^n$, $PS(\succ) \in \mathcal{A}$ is a random assignment where $A_{i,j}$ is the probability (fraction) that object j is assigned to (or “eaten by”) agent i .

2.2 General Properties

In this section we define key properties for matching mechanisms. To evaluate the quality of a random assignment, we use first-order stochastic dominance [14, 22]. Given a random assignment A_i , the probability that agent i is assigned an object that is at least as good as object ℓ is defined as follows

$$w(\succ_i, \ell, A_i) = \sum_{j \in M: j \succeq_i \ell} A_{i,j} \quad (1)$$

We say an agent always prefers assignment A_i to B_i , if for each object ℓ the probability of assigning an object at least as good as ℓ under A_i is greater or equal that of B_i , and strictly greater for some object.

Definition 1 (Stochastic Dominance) Given a preference ordering \succ_i , random assignment A_i stochastically dominates (sd) assignment $B_i (\neq A_i)$ if

$$\forall \ell \in M, w(\succ_i, \ell, A_i) \geq w(\succ_i, \ell, B_i) \quad (2)$$

A matching mechanism is *sd*-efficient if at all preference profiles $\succ \in \mathcal{P}^n$, for all agents $i \in N$, the prescribed assignment is not stochastically dominated by any other assignment.

Definition 2 (*sd*-Efficiency) A random assignment is *sd*-efficient if for all agents, it is not stochastically dominated by any other random assignment.

An important desirable property in matching mechanisms is strategyproofness, that is the mechanism is designed so that no agent has incentive to misreport its preferences.

² When $n < m$ and agents can receive more than one object, RSD requires a careful method for the picking sequence at each priority ordering to ensure strategyproofness. This picking sequence should be based on an arbitrary serial dictatorship quota mechanism, which directly affects the efficiency and envy of the assignments [15, 23]. For simplicity, we use the variant of RSD based on a *quasi-dictatorial* mechanism [40] where the first agent selects its most preferred $(m - n + 1)$ objects, and the rest of the agents choose one object each.

Definition 3 (*sd-Strategyproofness*) Mechanism \mathcal{M} is *sd-strategyproof* if at all preference profiles $\succ \in \mathcal{P}^n$, for all agents $i \in N$, and for any misreport $\succ'_i \in \mathcal{P}^n$, such that $A = \mathcal{M}(\succ)$ and $A' = \mathcal{M}(\succ'_i, \succ_{-i})$, we have:

$$\forall \ell \in M, w(\succ_i, \ell, A_i) \geq w(\succ_i, \ell, A'_i) \quad (3)$$

Sd-strategyproofness is a strict requirement. It implies that under any utility model consistent with the preference orderings, no agent can improve her expected utility by misreporting. We say that a mechanism is *weakly sd-strategyproof* if at all preference profiles there is no misreport such that for all $\ell \in M$, $w(\succ_i, \ell, A'_i) \geq w(\succ_i, \ell, A_i)$ with at least one $\ell' \in M$ such that $w(\succ_i, \ell', A'_i) > w(\succ_i, \ell', A_i)$. Clearly, *sd-strategyproofness* implies weak *sd-strategyproofness* but the converse does not hold.

An assignment is *manipulable* if it is not *sd-strategyproof*. If there exists some agent who strictly benefits from the manipulation, (*i.e.* the mechanism is not even weakly *sd-strategyproof*) then we say the assignment is *sd-manipulable* (or strictly manipulable).

We are also interested in whether mechanisms are fair and use the notion of *envyfreeness* to this end. An assignment is *sd-envyfree* if each agent strictly prefers her random allocation to any other agent's assignment.

Definition 4 (*sd-Envyfreeness*) Given agent i 's preference \succ_i , assignment A_i is *sd-envyfree* if for all agents $\forall k \neq i \in N$,

$$\forall \ell \in M, w(\succ_i, \ell, A_i) \geq w(\succ_i, \ell, A_k) \quad (4)$$

We say an assignment is weakly *sd-envyfree* if the inequality in Equation 4 is strict for some $\ell \in M$, but there exists at least one ℓ' for which the inequality in Equation 4 does not hold. A matching mechanism satisfies *sd-envyfreeness* if at all preference profiles $\succ \in \mathcal{P}^n$, it induces *sd-envyfree* assignments for all agents.

Lastly, we are interested in investigating efficiency, manipulation, and envy of the random mechanisms when preferences are lexicographic. Under lexicographic preferences, given two allocations, an agent prefers the one in which there is a higher probability for getting the most-preferred object.

Definition 5 (*Lexicographic Dominance*) Given a preference ordering \succ_i , random assignment A_i lexicographically dominates (*ld*) assignment B_i if

$$\begin{aligned} &\exists \ell \in M : w(\succ_i, \ell, A_i) > w(\succ_i, \ell, B_i) \wedge \\ &\forall k \succ_i \ell : w(\succ_i, \ell, A_i) = w(\succ_i, \ell, B_i). \end{aligned} \quad (5)$$

We say that allocation A *lexicographically dominates* another allocation B if there exists no agent $i \in N$ that lexicographically prefers B_i to A_i . Thus, an allocation mechanism is lexicographically efficient (*ld-efficient*) if for all preference profiles its induced allocation is not lexicographically dominated by any other random allocation.

Table 1: Properties of PS and RSD.

	$n \geq m$		$n < m$	
	PS	RSD	PS	RSD
<i>sd</i> -strategyproof	weak	✓	✗	✓
<i>sd</i> -efficiency	✓	✗	✓	✗
<i>sd</i> -envyfree	✓	weak	✓	weak

Table 2: Example showing the inefficiency of RSD

(a) Assignment under $PS(\succ)$					(b) Assignment under $RSD(\succ)$				
	a	b	c	d		a	b	c	d
A_1	1/2	0	1/2	0	A_1	5/12	1/12	5/12	1/12
A_2	1/2	0	1/2	0	A_2	5/12	1/12	5/12	1/12
A_3	0	1/2	0	1/2	A_3	1/12	5/12	1/12	5/12
A_4	0	1/2	0	1/2	A_4	1/12	5/12	1/12	5/12

2.3 Properties of RSD and PS

The theoretical properties of PS and RSD have been well studied in the economics literature [14], and we summarize the results in Table 1. Both mechanisms are ex post efficient, that is, their *realized outcomes* cannot be improved without making at least one agent worse off. PS has been shown to be both *sd*-envyfree and *sd*-efficient. However, it is not even weakly *sd*-strategyproof when $n < m$ [28] and is only weakly *sd*-strategyproof when $n \geq m$. On the other hand, RSD is always *sd*-strategyproof, but it is only weakly *sd*-envyfree and is not *sd*-efficient. Example 1 illustrates the *sd*-inefficiency of RSD.

Example 1 Suppose there are four agents $N = \{1, 2, 3, 4\}$ and four objects $M = \{a, b, c, d\}$. Consider the following preference profile $\succ = ((abcd), (abcd), (badc), (badc))$. Table 2 shows the outcomes for $PS(\succ)$ and $RSD(\succ)$. In this example, all agents strictly prefer the assignment induced by PS over the RSD assignment. Thus, RSD is inefficient at this preference profile.

3 Incomparability of RSD and PS

We argue that the theoretical findings on RSD and PS do not necessarily provide enough guidance to a market designer trying to select the correct mechanism for a specific setting. For example, while we know that PS is *sd*-efficient and RSD is not, this does not mean that PS assignment always stochastically dominate the assignments prescribed by RSD.

Example 2 Suppose there are three agents $N = \{1, 2, 3\}$ and three objects $M = \{a, b, c\}$. Consider the following preference profile $\succ = ((acb), (abc), (bac))$. Table 3 shows $PS(\succ)$ and $RSD(\succ)$. Neither assignment dominates the other since agent 1 is ambivalent between the two assignments while agent 2 prefers $PS(\succ)$ and agent 3 prefers $RSD(\succ)$.

Table 3: Incomparability of RSD and PS

(a) Assignment under $PS(\succ)$

	a	b	c
A_1	1/2	0	1/2
A_2	1/2	1/4	1/4
A_3	0	3/4	1/4

(b) Assignment under $RSD(\succ)$

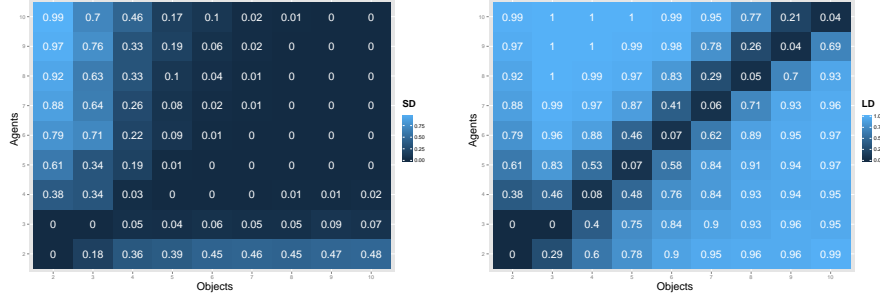
	a	b	c
A_1	1/2	0	1/2
A_2	1/2	1/6	1/3
A_3	0	5/6	1/6

If we knew the utility functions of the agents, consistent with their ordinal preferences, then we might be able to use the notion of (utilitarian) social welfare to help determine the better assignment.³ However, it is easy to construct different utility functions for the agents in Example 2 where both RSD and PS maximize social welfare. Similarly, the envy of RSD and the manipulability of PS both depend on the structure of preference profiles, and thus, a compelling question, that justifies studying the practical implications of deploying a matching mechanism, is to analyze the space of preference profiles to find the likelihood of inefficient, manipulable, or envious assignments under these mechanisms. In Example 2, for instance, if utilities of agents 1 and 2 are 10, 9, and 1, and agent 3's utility is 10, 6, and 4 for the first, second, and third objects respectively, then PS assignment outperforms that of RSD with respect to social welfare because $(\frac{1}{2} \cdot 10 + \frac{1}{2} \cdot 9 + (0) \cdot 1) + (\frac{1}{2} \cdot 10 + \frac{1}{4} \cdot 9 + \frac{1}{4} \cdot 1) + (\frac{3}{4} \cdot 10 + (0) \cdot 6 + \frac{1}{4} \cdot 4) > (\frac{1}{2} \cdot 10 + \frac{1}{2} \cdot 9 + (0) \cdot 1) + (\frac{1}{2} \cdot 10 + \frac{1}{6} \cdot 9 + \frac{1}{3} \cdot 1) + (\frac{5}{6} \cdot 10 + (0) \cdot 6 + \frac{1}{6} \cdot 4)$. However, if utility functions change such that all agents have the same utilities of 10, 9, and 1 for the first, second, and third objects respectively, then the social welfare under RSD outperforms that of PS because $(\frac{1}{2} \cdot 10 + \frac{1}{2} \cdot 9 + (0) \cdot 1) + (\frac{1}{2} \cdot 10 + \frac{1}{6} \cdot 9 + \frac{1}{3} \cdot 1) + (\frac{5}{6} \cdot 10 + (0) \cdot 9 + \frac{1}{6} \cdot 1) > (\frac{1}{2} \cdot 10 + \frac{1}{2} \cdot 9 + (0) \cdot 1) + (\frac{1}{2} \cdot 10 + \frac{1}{4} \cdot 9 + \frac{1}{4} \cdot 1) + (\frac{3}{4} \cdot 10 + (0) \cdot 9 + \frac{1}{4} \cdot 1)$.

4 General and Lexicographic Preferences

The theoretical properties of PS and RSD only provide limited insight into their practical applications. In particular, when deciding which mechanism to use in different settings, the incomparability of PS and RSD leaves us with an ambiguous choice in terms of efficiency, manipulability, and envyfreeness. Thus, we examine the properties of RSD and PS in the space of all possible preference profiles as well as under lexicographic preferences. Lexicographic preferences are present in various applications and have been extensively studied in artificial intelligence and multiagent systems as a means of assessing allocations based on ordinal preferences [18, 21, 44]. Under lexicographic preferences, an allocation that assigns a higher probability to the top ranked object is always preferred to any other allocation, regardless of the probabilities assigned to objects in the next positions. Only when two allocations assign equal probabilities to the top ranked object, the probability of the next preferred object is considered. In the rest of this paper, we denote the efficiency, strategyproofness, manipulability, and envyfreeness with ld - (lexicographically dominate) prefix.

³ Given utility functions for the agents, where $u_i(j)$ is the utility agent i derives from being assigned object j , the (utilitarian) social welfare of an assignment A is $\sum_i \sum_j A_{i,j} u_i(j)$.



(a) The fraction that PS stochastically dominates RSD. (b) The fraction that PS lexicographically dominates RSD.

Fig. 1: The fraction of preference profiles under which PS dominates RSD.


The number of all possible preference profiles is super exponential $(m!)^n$. For each combination of n agents and m objects we performed a brute force coverage of all possible preference profiles. Thus, for all subsequent figures each data point shows the fraction of all possible preference profiles. For the cases of $n = 10$ and $m \in \{9, 10\}$, we randomly generated 1,000 instances by sampling from a uniform preference profile distribution. For each preference profile, we ran both PS and RSD mechanisms and compared their outcomes in terms of the stochastic dominance relation. Appendix A illustrates our numerical results. Note that not only is computing RSD probabilities $\#P$ -complete (and thus intractable) [6, 45], but checking the desired properties such as envyfreeness, efficiency, and manipulability of random allocations is shown to be NP-hard for general settings [9, 10]. Thus, for larger settings even if we randomly sample preference profiles it is not easy to verify the aforementioned properties.

4.1 Preliminary Results

Our experimentation disclose several intriguing observations, confirming theoretical results and providing additional insights into matching markets. A preliminary look at our empirical results illustrates the following: when $m \leq 2, n \leq 3$, PS coincides exactly with RSD, which results in the best of the two mechanisms, *i.e.*, both mechanisms are *sd*-efficient, *sd*-strategyproof, and *sd*-envyfree. Another interesting observation is that when $m = 2$, for all n , PS is *sd*-strategyproof (although the PS assignments are not necessarily equivalent to assignments induced by RSD), RSD is *sd*-envyfree, and for most instances when $m = 2$, PS stochastically dominates RSD, particularly when $n \geq 4$.

4.2 Efficiency

Our first finding is that the fraction of preference profiles at which RSD and PS prescribe identical random assignments goes to 0 when n grows. There are two conclusions that one can draw. First, this result confirms the theoretical results



(a) The fraction of manipulable preference profiles under PS. (b) The fraction of *sd*-manipulable profiles under PS.

Fig. 2: Heatmaps illustrating the manipulability of PS.

of Manea on asymptotic inefficiency of RSD [33], in that, in most instances, the assignments induced by RSD are not identical to the PS assignments. Second, this result suggests that the incomparability of outcomes is significant, that is, the social welfare of the random outcomes is highly dependent on the underlying utility models.

The fraction of preference profiles $\succ \in \mathcal{P}^n$ for which RSD is stochastically dominated by PS at \succ converges to zero as $\frac{n}{m} \rightarrow 1$. Figure 1a shows that when m grows beyond $m > 5$, due to incomparability of RSD and PS with regard to the stochastic dominance relation, the RSD assignments are rarely stochastically dominated by *sd*-efficient assignments prescribed by PS.

We also see similar results when we restrict ourselves to lexicographic preferences (Figure 1b). The fraction of preference profiles $\succ \in \mathcal{P}^n$ for which RSD is lexicographically dominated by PS at \succ converges to zero as $\frac{n}{m} \rightarrow 1$.

For lexicographic preferences, we also observe that the fraction of preference profiles for which PS assignments strictly dominate RSD-induced allocations goes to 1 when the number of agents and objects diverge. The fraction of preference profiles $\succ \in \mathcal{P}^n$ for which RSD is lexicographically dominated by PS at \succ converges to 1 as $|n - m|$ grows. Intuitively, when some agents can receive more than one object ($n < m$) or when there are not sufficient objects ($n > m$) for all agents, in each realized ordering of agents by RSD, those with higher priority are treated very differently than those in lower priority. Thus, the RSD outcome tend to be unfair and undesirable for most agents.

One immediate conclusion is that although RSD does not guarantee either *sd*-efficiency or *ld*-efficiency, in most settings when $\frac{n}{m} \rightarrow 1$ (and also $n \leq m$ for *sd*-efficiency according to Figure 1a), neither of the two mechanisms is preferred in terms of efficiency. Hence, one cannot simply rule out the RSD mechanism.

4.3 Manipulability of PS

One critical issue with deploying PS is that it does not provide incentives for honest reporting of preferences. Although for $n \geq m$ PS is weakly *sd*-strategyproof [14]

and ld -strategyproof [46], when $n < m$ PS no longer satisfies these two properties.⁴ The real concern is that, in the absence of strategyproofness, PS allocations are only efficient (or envyfree) with respect to the reported preferences. Thus, if an agent decides to manipulate the outcome by misreporting its preferences, PS will no longer guarantee efficiency, nor envyfreeness with respect to the true underlying preferences. Thus, we are interested in understanding the degree to which PS allocations are manipulable.

Figure 2 shows that the fraction of manipulable profiles goes to 1 as n or m grow. PS is almost 99% manipulable for $n > 5, m > 5$. Another interesting observation is that, for all $n < m$, the fraction of sd -manipulable preference profiles goes to 1 as $m - n$ grows (Figure 2b). These results imply that when agents are permitted to receive more than a single object, agents can strictly benefit from misreporting their preferences.

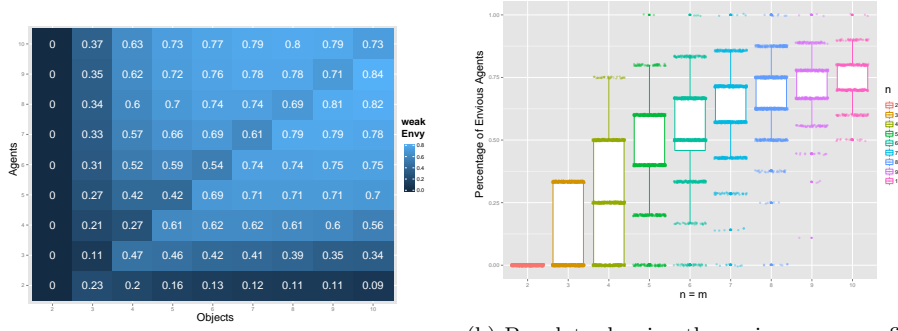
Moreover, at those instances of problem where PS is sd -strategyproof, the assignment prescribed by PS most often coincides with the RSD induced assignment. For example, when $n = m = 5$, PS is only sd -strategyproof at 11% of preference profiles, 7% of which are identical to the assignments induced by RSD. This insight further confirms the vulnerability of PS to misreporting (See Table 6 for detailed numerical results).



Fig. 3: The fraction of ld -manipulable profiles under PS.

As illustrated in Figure 3, the manipulability of PS under lexicographic preferences has a similar trend when there are more objects than agents ($n < m$) and the fraction of ld -manipulable preference profiles converges to 1 even more rapidly when $\frac{m}{n}$ grows.

⁴ A recent experimental study on the incentive properties of PS shows that human subjects are less likely to manipulate the mechanism when misreporting is a Nash equilibrium. However, subjects' tendency for misreporting is still significant even when it does not improve their allocations [24].



(a) A heatmap showing the percentage of envious agents.

(b) Boxplots showing the various envy profiles for $n = m$. The Y axis represents the percentage of envious agents.

Fig. 4: Plots representing the percentage of (weakly) envious agents under RSD.

4.4 Envy in RSD

The PS mechanism has a desirable fairness property and is guaranteed to satisfy *sd*-envyfreeness, whereas RSD is not *sd*-envyfree. To further investigate the envy among agents under RSD, we measured the fraction of agents that are weakly *sd*-envious of at least one other agent.

Figure 4 shows that for RSD, the percentage of agents that are weakly envious increases with the number of agents. Figure 4a reveals an interesting observation: fixing any $n > 3$, the percentage of agents that are (weakly) envious grows with the number of objects, however, there is a sudden drop in the percentage of envious agents when there are equal number of agents and objects.

For better understanding of the population of agents who feel (weakly) envious under RSD, we illustrate the various envy profiles based on the percentage of envious agents in all instances of the problem when $n = m$ (Figure 4b). One observation is that there are few distinct envy profiles at each n , each representing a particular class of preference profiles, and by increasing n , the fraction of agents that are envious of at least one agent increases.

5 Utility Models

Given a utility model consistent with an agent's preference ordering, we can find the agent's expected utility for a random assignment. Let u_i denote agent i 's Von Neumann-Morgenstern (VNM) utility model consistent with its preference ordering \succ_i . That is, $u_i(a) > u_i(b)$ if and only if $a \succ_i b$. Then, agent i 's expected utility for random assignment A_i is $\mathbb{E}(u_i|A_i) = \sum_{j \in M} A_{i,j} u_i(j)$.

We say that agent i (strictly) prefers assignment A_i to B_i if and only if $\mathbb{E}(u_i|A_i) > \mathbb{E}(u_i|B_i)$. A mechanism is strategyproof if there exists no agent that can improve its expected utility by misreporting its preference ordering.

Definition 6 (Strategyproof) Mechanism \mathcal{M} is strategyproof if for all agents $i \in N$, and for any misreport $\succ'_i \in \mathcal{P}^n$, such that $A = \mathcal{M}(\succ)$ and $A' = \mathcal{M}(\succ'_i, \succ_{-i})$, given a utility model u_i consistent with \succ_i , we have $\mathbb{E}(u_i|A_i) \geq \mathbb{E}(u_i|A'_i)$.

A matching mechanism is envyfree if for all preference profiles it prescribes an envyfree assignment.

Definition 7 (Envyfreeness) Assignment A is envyfree if for all $i, k \in N$, given utility model u_i consistent with \succ_i , we have $\mathbb{E}(u_i|A_i) \geq \mathbb{E}(u_i|A_k)$.

Given utility functions for the agents, the (utilitarian) *social welfare* of an assignment A is $\sum_i \mathbb{E}(u_i|A_i)$. A random assignment A is *sd-efficient* if and only if there exists a profile of utility values consistent with \succ such that A maximizes the social welfare ex ante [14, 37]. This existence result does not shed light on the social welfare when comparing two random assignments, since an assignment can be *sd-efficient* but may not have desirable ex ante social welfare. Consider the following random assignments: assignment A which is *sd-efficient* and assignment $B \neq A$ which is not stochastically dominated by A . Given a preference profile, A is guaranteed to maximize the social welfare for at least one profile of consistent utilities. However, there may be other profiles of utilities consistent with preferences at which B maximizes the sum of utilities (social welfare).

Example 3 Consider the problem introduced in Example 2 with assignments illustrated in Table 3. Let's assume that all agents have the same utility model $u_1 = u_2 = u_3$ where the utilities are 10, 9, 0 for the first, second, and third objects respectively. The sum of expected utilities under the PS assignment is $(\frac{1}{2} \cdot 10 + \frac{1}{2} \cdot 9 + 0) + (\frac{1}{2} \cdot 10 + \frac{1}{4} \cdot 9 + \frac{1}{4} \cdot 0) + (\frac{3}{4} \cdot 10 + 0 \cdot 9 + \frac{1}{4} \cdot 0)$, while the sum of expected utilities under the RSD allocation is $(\frac{1}{2} \cdot 10 + \frac{1}{2} \cdot 9 + 0) + (\frac{1}{2} \cdot 10 + \frac{1}{6} \cdot 9 + \frac{1}{3} \cdot 0) + (\frac{5}{6} \cdot 10 + 0 \cdot 9 + \frac{1}{6} \cdot 0)$. It is easy to see that for this profile, the ex ante social welfare under RSD is larger than that of PS.

Thus, given a profile of utilities we investigate the (ex ante) social welfare of the assignments under PS and RSD.

5.1 Instantiating Utility Functions

To deepen our understanding as to the performance of the two mechanisms, we investigate different utility models. In particular we look at the performance of the mechanisms when the agents are all risk neutral (*i.e.* have linear utility functions), when agents are risk seeking and when agents are risk averse.

Our first utility model is the well-studied linear utility model. Given an agent i 's preference ordering \succ_i , we let $r(\succ_i, j)$ denote the rank of object j . For example, given preference ordering $a \succ_i b \succ_i c$ then $r(\succ_i, a) = 1$, $r(\succ_i, b) = 2$ and $r(\succ_i, c) = 3$. The utility function for agent i , given object j is $u_i(j) = m - r(\succ_i, j)$.

We use an *exponential* utility model to capture risk attitudes beyond risk-neutrality. An exponential utility has been shown to provide an appropriate translation for individuals' utility models [4]. In particular, we define the exponential utility as follows:

$$u_i(j) = \begin{cases} (1 - e^{-\alpha(m - r(\succ_i, j))})/\alpha, & \alpha \neq 0 \\ m - r(\succ_i, j), & \alpha = 0 \end{cases} \quad (6)$$

The parameter α represents the agent's risk attitude. If $\alpha > 0$ then the agent is risk averse, while if $\alpha < 0$ then the agent is risk seeking. When $\alpha = 0$ then the




Fig. 5: Utility values for various α under risk taking, risk neutral, and risk averse models. There are eleven objects ranked from 1 to 11, with linear utilities from 0 (the last object) to 10 (the top choice). The trendlines fit exponential trends to the discrete alpha parameters.

Table 4: Sample utility values when there are 3 objects under different risk attitudes and risk intensities.

rank / α	$\alpha = -2$	$\alpha = -1$	$\alpha = 0$	$\alpha = 1$	$\alpha = 2$
1	26.799	6.389	2	0.865	0.491
2	3.195	1.718	1	0.632	0.432
3	0	0	0	0	0

agent is risk neutral and we have a linear utility model. The value $|\alpha|$ represents the intensity of the attitude. That is, given two agents with $\alpha_1 > \alpha_2 > 0$, we say that agent 1 is more risk averse than agent 2. Similarly if $\alpha_1 < \alpha_2 < 0$ then agent 1 is more risk seeking than agent 2. Figure 5 illustrates the risk curvature for various risk taking and risk averse α parameters.

Table 4 shows sample utility values for various risk taking, neutral, and risk averse utility profiles. These values show how a utility for objects in various ranking positions will change according to risk attitude models. Note that we do normalize the utilities such that all utilities add up to 1.

6 Results

For our experiments, we vary three parameters: the number of agents n , the number of objects m , and the risk attitude factor α . Each data point in the graphs shows the average over all possible preference profiles. We study the same settings as in Section 4 when $n \geq m$ and $n < m$. For each utility function, we look at homogeneous populations of agents where agents have the same risk attitudes but may have difference ordinal preferences.

To compare the social welfare, we investigate the percentage change (or improvement) in social welfare of PS compared to RSD under various utility models. That is,

$$\frac{\sum_i \mathbb{E}(u_i | PS(\succ)) - \sum_i \mathbb{E}(u_i | RSD(\succ))}{\sum_i \mathbb{E}(u_i | RSD(\succ))}.$$

To measure the manipulability of PS, we are interested in answering two key questions: *i*) In what fraction of profiles is PS manipulable by at least one agent? and *ii*) If manipulation is possible, what is the average percentage of maximum gain? That is,

$$\max_i \left\{ \frac{\mathbb{E}(u_i | PS(\succ'_i, \succ_{-i})) - \mathbb{E}(u_i | PS(\succ))}{\mathbb{E}(u_i | PS(\succ))} \right\}.$$

To study the envy under the RSD mechanism, we consider two measures: *i*) the fraction of envious agents, and *ii*) the total envy felt by all agents.

6.1 Linear Utility Model

We first looked at how RSD and PS perform under the assumption that the utility models are linear (Figure 6). In most cases, the social welfare under PS increases compared to RSD; however, the social welfare of PS is very close to that of RSD when $n = m$ (less than 0.015 overall improvement in all cases). Interestingly, under RSD the fraction of envious agents gets close to 0 when $n \geq m$. With regards to strategyproofness, PS is manipulable in most combinations of n and m and the fraction of manipulable profiles and the utility gain from manipulation increases as the number of objects compared to agents increases.

6.2 Risk Seeking

Figure 7 presents our results in terms of percentage change in social welfare between PS and RSD. Positive numbers show the percentage of improvement in social welfare from PS to RSD. Negative values represent those cases where RSD has increased social welfare compared to PS.

Social welfare: Fixing $\alpha < 0$, for $n \geq m$ when $\frac{n}{m}$ grows PS improves the social welfare compared to RSD in all instances of the problem and the percentage of improvement also increases. A similar trend holds when varying risk intensity α for fixed n and m where $n \neq m$. For $n < m$, when $\frac{m}{n}$ grows the fraction of profiles at which PS has higher social welfare compared to RSD rapidly increases and the percentage change is also noticeably larger, quickly getting close to 90% improvement (Figures 7a, 7c, and 7e). This social welfare gap between PS and RSD grows as the risk intensity $|\alpha|$ increases. Surprisingly, this trend changes for equal number of agents and objects $n = m$: the more risk-seeking agents are (larger $|\alpha|$), RSD becomes more desirable than PS, and in fact, RSD improves the social welfare in more instances.

Envy: Figure 8 shows that for $n \geq m$, the fraction of envious agents under all profiles vanishes and RSD becomes envyfree. This is more evident when agents are more risk-seeking. Intuitively, these observations confirm the theoretical findings about the envyfreeness of RSD under lexicographic preferences [23]. This is because one can consider lexicographic preferences as risk-seeking preferences where an object in a higher ranking is infinitely preferred to all objects that are ranked

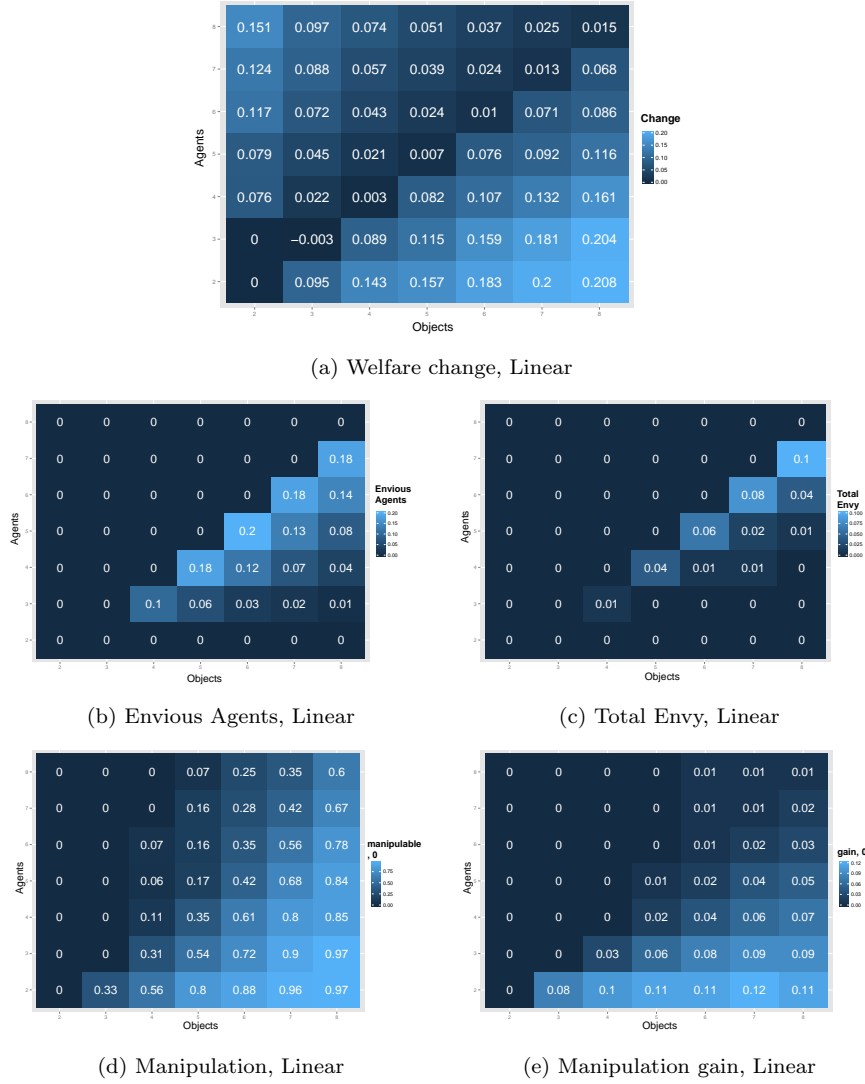


Fig. 6: Linear Utility: Welfare Change, Envy, and Manipulation.

less preferably [23]. When $n < m$, our quasi-dictatorial extension of RSD creates some envy among the agents, because the agent with the highest priority receives $m - n + 1$ objects, while all other agents receive at most one object. An interesting result is the envy created by RSD starts to fade out when the risk intensity $|\alpha|$ increases.

Manipulability: Figure 9 shows the manipulability of the PS assignments when agents are risk seeking. We see that the possibility of manipulation (and any gain) decreases as the risk intensity increases. When $n \geq m$ the fraction of manipulable profiles goes to 0 the more risk seeking agents become. However, when

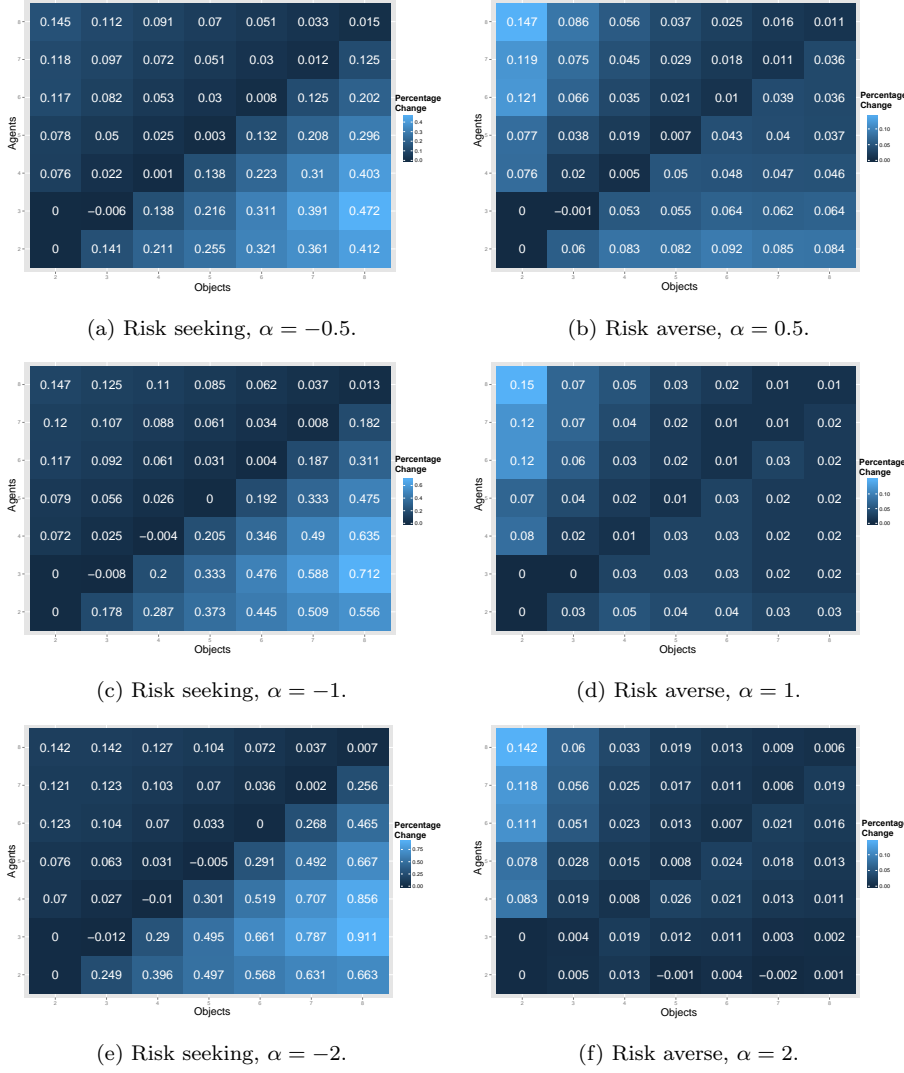


Fig. 7: The percentage change in social welfare. The negative values show that RSD outperforms PS.

$n < m$ even though the fraction of manipulable profiles (and manipulation gain) decreases, the fraction of manipulable profiles goes to 1 as $\frac{m}{n}$ grows.

6.3 Risk Aversion

Social welfare: Figures 7b, 7d, and 7f show that, fixing risk factor $\alpha > 0$, when $\frac{n}{m}$ grows, PS assignments are superior to that of RSD in terms of social welfare in more instances, and the percentage change in social welfare increases. Fixing risk




(a) Fraction of envious agents, $\alpha = -0.5$.(b) Total envy, $\alpha = -0.5$.(c) Fraction of envious agents, $\alpha = -1$.(d) Total envy, $\alpha = -1$.(e) Fraction of envious agents, $\alpha = -2$.(f) Total envy, $\alpha = -2$.

Fig. 8: The fraction of envious agents and total envy perceived by agents under risk-seeking utilities.

factor $\alpha > 0$ and when $\frac{m}{n}$ grows, RSD is more likely to have the same social welfare as PS, and in fact in some instances the social welfare under RSD is better than the social welfare under PS. Fixing m and n , when the risk intensity α increases RSD is more likely to have the same social welfare as PS, that is, the welfare gap between PS and RSD closes when agents are more risk averse (α increases). This result is insightful and states that under risk aversion the random allocations prescribed by RSD are either as good as PS or in some cases even are superior to the allocations prescribed by PS due to the underlying shape of the utility

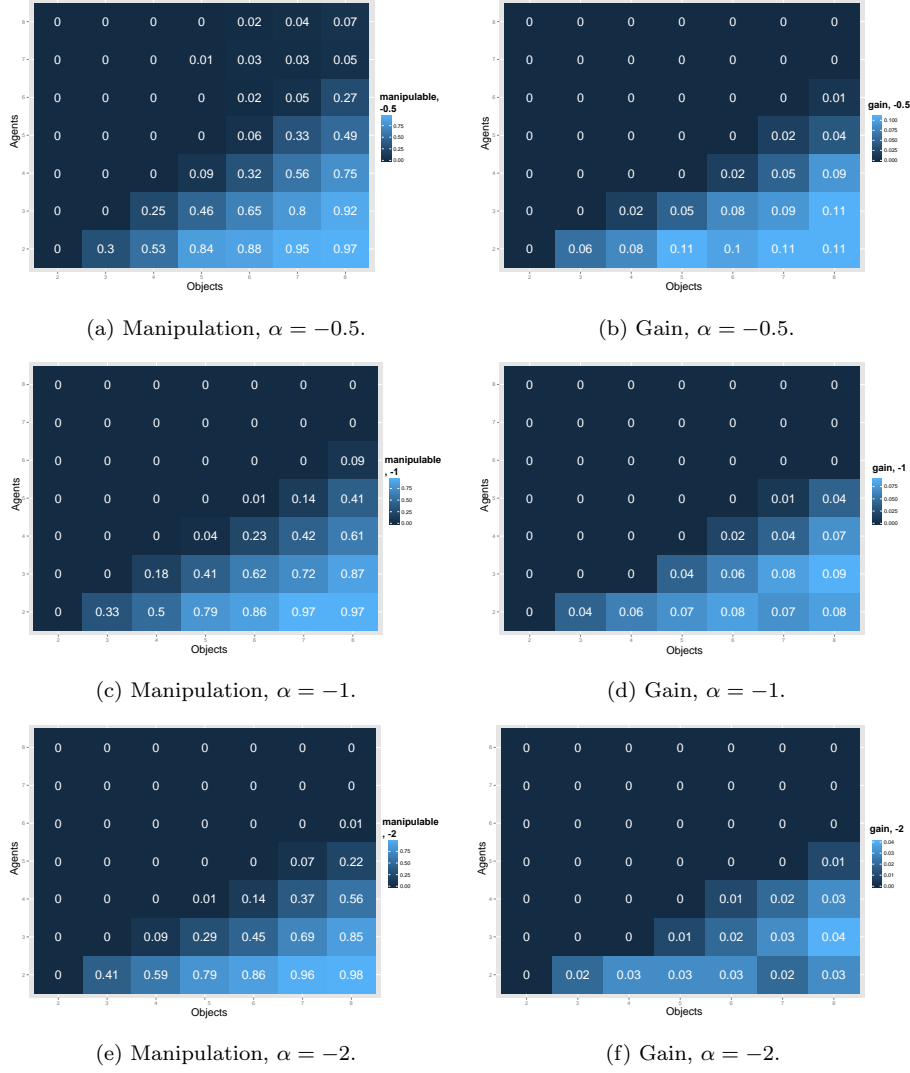


Fig. 9: The fraction of manipulable instances and manipulation gain of PS under risk-seeking preferences.

models. Figure 17 illustrates the percentage change in social welfare based on the difference between available objects and agents ($m - n$) for risk seeking, linear, and risk averse utilities with different risk intensities.

Envy: In Figure 10 we observe that when $n \geq m$, the fraction of envious agents and total envy grows as $\frac{n}{m} \rightarrow 1$. Increasing the risk intensity ($|\alpha|$), the fraction of envious agents increases; however, the total envy among the agents remains considerably low. For $n < m$, the fraction of envious agents and total envy grows as risk intensity increases. An interesting observation is that envy is maximized




(a) Fraction of envious agents, $\alpha = 0.5$.(b) Total envy, $\alpha = 0.5$.(c) Fraction of envious agents, $\alpha = 1$.(d) Total envy, $\alpha = 1$.(e) Fraction of envious agents, $\alpha = 2$.(f) Total envy, $\alpha = 2$.

Fig. 10: The fraction of envious agents and total envy perceived by agents under risk aversion. The total envy is shown up to two decimal points.

when $m = n + 1$, and it decreases as $\frac{m}{n}$ grows. This is mostly due to the choice of using randomized quasi-dictatorial mechanism for implementing RSD where the first dictator receives $m + n - 1$ objects and all other agents only receive a single object. Lastly, we noticed that in all instance where RSD creates envy among the agents, around 25% of agents bear more than 50% of envy. That is, few agents feel extremely envious while all other agents are either envyfree or only feel a minimal amount of envy.

Manipulability: Figures 11 illustrates the manipulability of the PS assignments when agents have risk averse preferences. The fraction of manipulable pro-

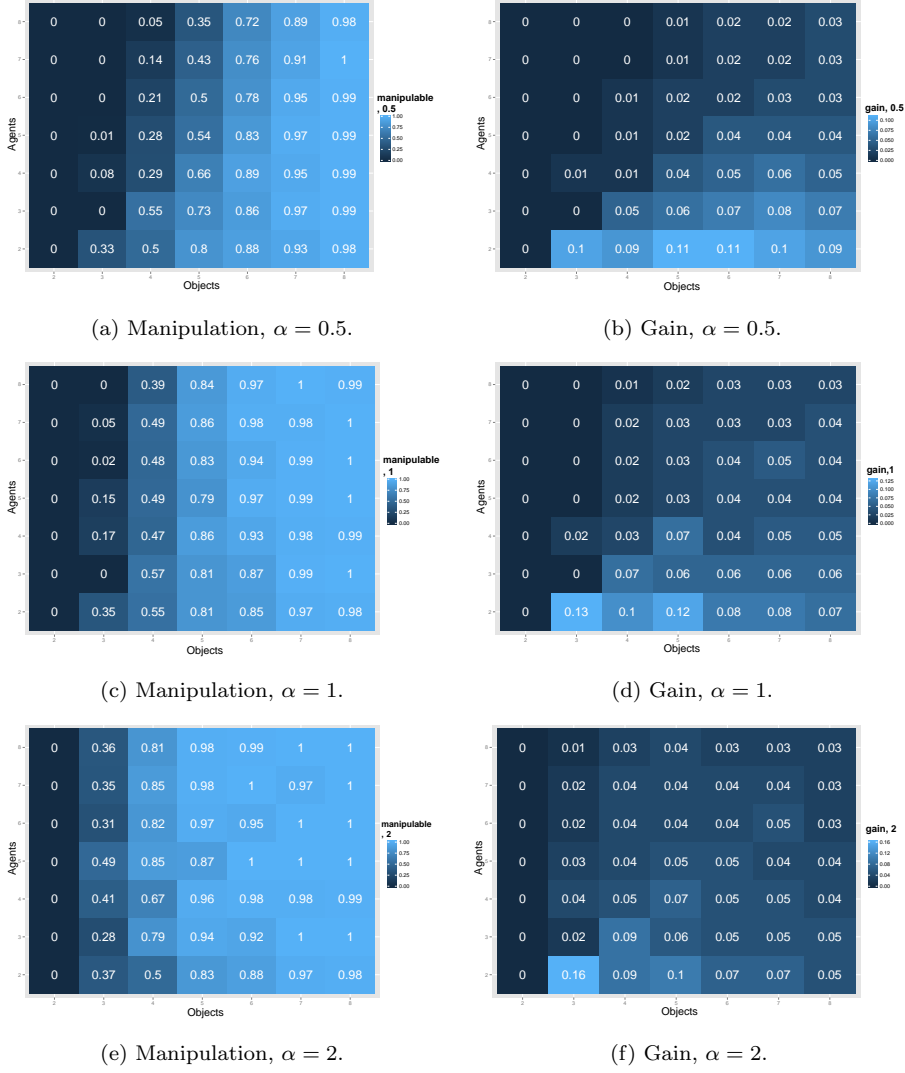


Fig. 11: The fraction of manipulable instances and manipulation gain of PS under risk aversion.

files rapidly goes to 1 as $\frac{m}{n}$ grows. Similarly, as agents become more risk averse (α increases) the fraction of manipulable profiles goes to 1 and the manipulation gain increases.

7 Other Ranking Distribution Models

In this section, we use variations of two statistical models that are commonly used to capture realistic preference distributions in a population of players. Considerable

work in computational social choice and machine learning has exploited these statistical models to capture the distribution of ranking preferences in a population of agents [8, 30, 31]. We will focus on Mallows Models and Polya-Eggenberger Urn Models (Urn) [11, 32, 35].

In Mallows models the population is distributed around a reference ranking proportional to the Kendall-Tau (KT) distance [25, 26]. Henceforth, preferences closer to the reference ranking are more likely to appear in the population. In other words, agents' preferences deviate from the reference ranking with decreasing probability as rankings move away from the reference. Mallows models are parametrized by a reference ranking and a dispersion parameter. Formally, given a reference ranking (\succ) and a dispersion parameter (ϕ), we have

$$P(\succ) = \frac{1}{Z} \phi^{KT(\succ, \hat{\succ})}, \quad \forall \succ \in \mathcal{P} \quad (7)$$

where $Z = 1 \cdot (1 + \phi) \cdot (1 + \phi + \phi^2) \dots (1 + \dots + \phi^{m-1})$. When $\phi = 1$ the Mallows model is equivalent to the uniform distribution, and when $\phi = 0$ the distribution mass is entirely on the reference ranking. It is also possible for an agent population to have multiple references. In these cases, Mallows Mixture models are parametrized by a set of ranking reference with their corresponding dispersion parameters.

In the Urn distribution model, with every random selection of a preference order the probability of this preference order being selected in subsequent samples increases. Intuitively, we can think of a collection of $m!$ preference orderings and every time an ordering is sampled uniformly from this collection, it will be replaced by two copies of the same preference ordering.

We use the PrefLib Toolkit [36] to generate Mallows and Urn distribution models. In our experiments, we used Mallows model with one reference ranking as well as Mallows mixture models with five reference rankings. Every data point in the figures is averaged over 1,000 samples.

In general, the same patterns under the uniform preference distribution hold for various numbers of agents and objects and when varying the risk parameter and utility functions.

Social welfare: Figures 12 and 13 show the results of our simulation for social welfare when agents' preferences are drawn from Mallows models. These results are consistent with pure Mallows models with single reference rankings (Figure 12) as well as Mallows mixture models with five references (Figure 13). The percentage change in social welfare is infinitesimal when $n \geq m$ for both risk averse and risk seeking populations. Under risk aversion, this percentage change in social welfare remains small when $n < m$. Similar to the uniform populations, the negative values show that in some cases, particularly when $n = m$, RSD assignments outperform those under PS assignments.

Manipulability: The PS assignments remain very susceptible to manipulation even under more natural assumptions on how the preferences are distributed. Figures 14 illustrates the manipulability of the PS assignments and the average gain from manipulation when agents are drawn from Mallows mixture models under risk averse and risk seeking attitudes.

The fraction of manipulable profiles and manipulation gain goes to 0 when agents are risk seeking. Under risk aversion, the fraction of manipulable profiles and manipulation gain rapidly increases as $\frac{m}{n}$ grows (Figure 14). These results hold




Fig. 12: The percentage change in social welfare under the pure Mallows distribution with single reference.

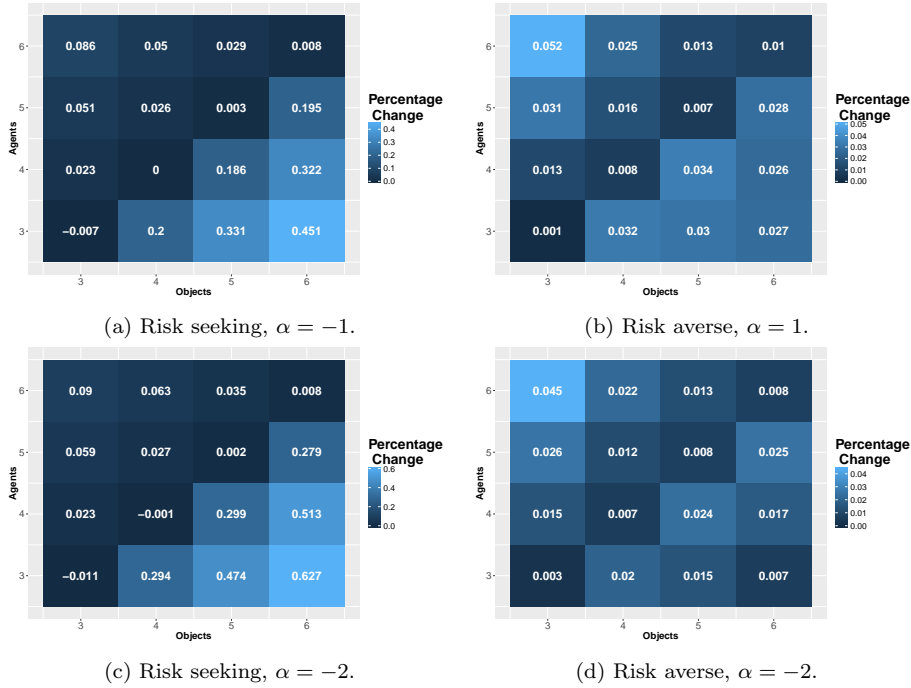


Fig. 13: The percentage change in social welfare under the Mallows Mixture distribution with five references.

under pure Mallows distribution (Figure 15) as well as under the Polya-Urn model (Figure 16) but with slightly slower growth. This is consistent with the fact that, in less diverse populations, agents' preference are more similar and conflicting, and thus manipulation is less likely (even though still significantly considerable) as opposed to more diverse and uniform set of profiles.

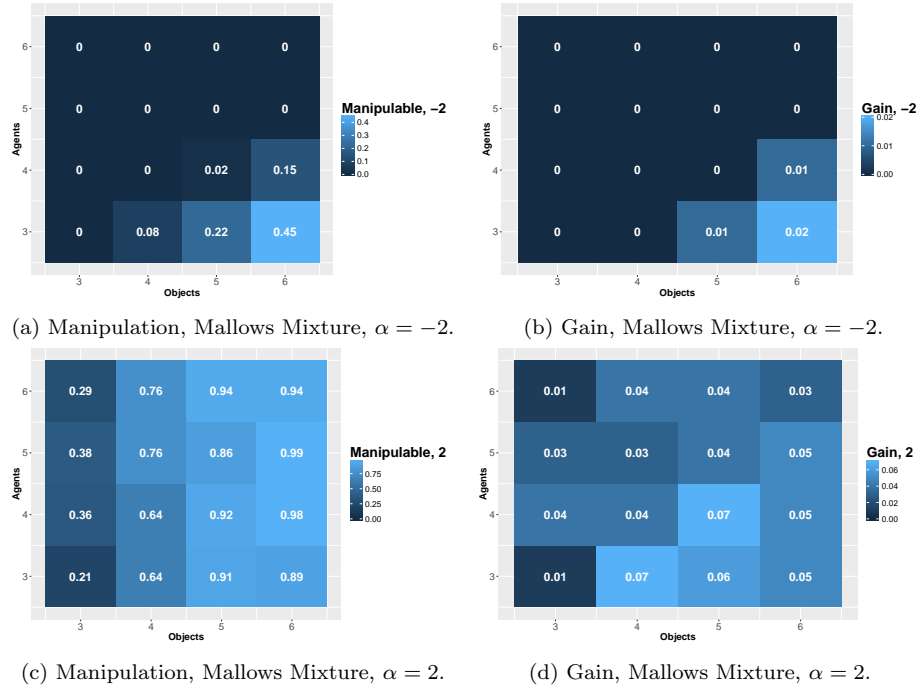


Fig. 14: The fraction of manipulable instances and manipulation gain of PS under the Mallows mixture model with five references.

8 Related Literature

Assignment problems with ordinal preferences have attracted interest from many researchers. Svensson showed that serial dictatorship is the only deterministic mechanism that is strategyproof, nonbossy, and neutral [48]. Random Serial Dictatorship (RSD) (uniform randomization over all serial dictatorship assignments) satisfies strategyproofness, proportionality, and ex post efficiency [2]. Bogomolnaia and Moulin noted the inefficiency of RSD from the ex ante perspective, and characterized the matching mechanisms based on first-order stochastic dominance [14]. They proposed the probabilistic serial mechanism as an efficient and envyfree mechanism with regards to ordinal preferences. While PS is not strategyproof, it satisfies weak strategyproofness for problems with equal number of agents and objects. However, PS is strictly manipulable (not weakly strategyproof) when there are more objects than agents [27]. Kojima and Manea, showed that in large assignment problems with sufficiently many copies of each object, truth-telling is a weakly dominant strategy in PS [28]. In fact PS and RSD mechanisms become equivalent [17], that is, the inefficiency of RSD and manipulability of PS vanishes when the number of copies of each object approaches infinity.

The practical implications of deploying RSD and PS have been the center of attention in many one-sided matching problems [1, 38]. In the school choice setting with multi-capacity alternatives, Pathak observed that many students obtained a more desirable random assignment through PS in public schools of New York

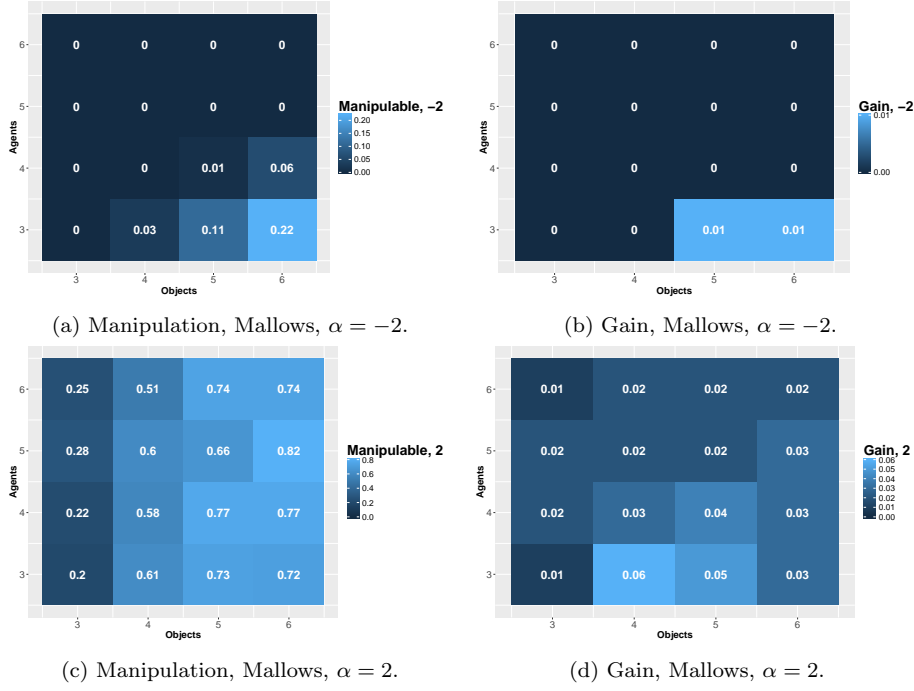


Fig. 15: The fraction of manipulable instances and manipulation gain of PS under the pure Mallows model with one reference ranking.

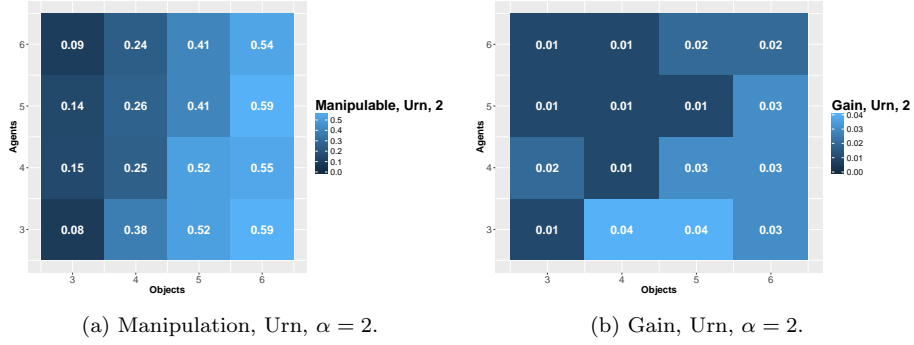


Fig. 16: The fraction of manipulable instances and manipulation gain of PS under the Polya-Urn model.

City [41]; however, the efficiency difference was quite small. These equivalence results and their extensions to all random mechanisms [29], do not hold when the quantities of each object is limited to one.

Other interesting aspects of PS and RSD such as computational complexity and best-response strategies have also been explored [8, 9, 19]. In this vein, Aziz et al. proved the existence of pure Nash equilibria, but showed that computing

Table 5: A random assignment for a preference profile wherein PS and RSD both prescribe an identical matching, i.e. $PS(\succ) = RSD(\succ)$.

	a	b	c
A_1	1/3	1/2	1/6
A_2	1/3	0	2/3
A_3	1/3	1/2	1/6

an equilibrium is NP-hard [8]. Nevertheless, Mennle et al. [39] showed that agents can easily find near-optimal strategies by simple local and greedy search. In the absence of truthful incentives, the outcome of PS is no longer guaranteed to be efficient or envyfree with respect to agents' true underlying preferences, and this inefficiency may result in outcomes that are worse than RSD, especially in 'small' markets [19]. The utilitarian and egalitarian welfare guarantees of RSD have been studied under ordinal and linear utility assumptions [7, 12]. For arbitrary utilities, RSD provides the best approximation ratio for utilitarian social welfare when $m = n$ among all mechanisms that rely only on ordinal preferences [20].

9 Discussion

We studied the space of general preferences and provided empirical results on the incomparability of RSD and PS. It is worth mentioning that at preference profiles where PS and RSD induce identical assignments, RSD is *sd*-efficient, *sd*-envyfree, and *sd*-strategyproof. However, PS is still highly manipulable. We further strengthen this argument by providing an observation in Example 4:

Example 4 Consider the following preference profile $\succ = ((bca), (cab), (bca))$. Table 5 shows the prescribed random assignment. In this example, with PS as the matching mechanism, agent 1 can misreport her preference as $\succ'_1 = (cba)$, and manipulate her assignment to $1/4(b), 1/2(c), 1/4(a)$. It is easy to see that agent 1's misreport improves her expected outcome for all utility models where $\frac{2}{6}u_1(c) > \frac{1}{4}u_1(b) + \frac{1}{12}u_1(a)$ (for example utilities 10, 9, 0 for b, c, a respectively.).

We investigated various utility models according to different risk attitudes. Our findings hold under various assumptions on the population of agents and preference profile distributions. Our main results are:

- In terms of efficiency, the fraction of preference profiles $\succ \in \mathcal{P}^n$ for which PS stochastically (or lexicographically) dominates RSD converges to zero as $\frac{n}{m} \rightarrow 1$. When instantiating the preferences with actual utility functions, PS allocations are only slightly better than RSD allocations in terms of social welfare when varying n and m , particularly under risk averse utilities. In fact, in some cases RSD allocations are superior in terms of social welfare (see Figure 17).
- PS is almost 99% manipulable when $n \leq m$ and the fraction of *sd*- and *ld*-manipulable profiles rapidly goes to 1 as $\frac{m}{n}$ grows. When instantiating the preferences with utility functions, the manipulability of PS increases as agents become more risk averse. Moreover, an agent's utility gain from manipulation also grows when the risk intensity increases.



Fig. 17: The percentage change in social welfare between RSD and PS for $\alpha \in (-2, -1, -0.5, 0, 0.5, 1, 2)$ and different combinations of $m - n$. Positive α indicates risk averse and negative α risk taking profiles. Linear utility is indicated by $\alpha = 0$. As agents become more risk averse the social welfare gap between RSD and PS closes.

- For risk seeking utilities, when $n \geq m$ the fraction of envious agents under all profiles vanishes and RSD becomes envyfree. For risk averse utilities, the fraction of envious agents increases as agents become more risk averse. However, the total amount of envy just slightly grows, and surprisingly, only few agents feel extremely envious while all other agents are either envyfree or only feel a minimal amount of envy.

An interesting future direction is to study egalitarian social welfare of the matching mechanisms in single and multi unit assignment problems as well as in the full preference domain. Another open direction is to provide a parametric analysis of the matching mechanisms according to the risk aversion factor.

10 Design Recommendations for Multiagent Systems Practitioners

Our work in this paper can be used to help guide designers of multiagent systems who need to solve allocation problems. If a designer strongly requires *sd*-efficiency then the theoretical results of PS indicate that it is better than RSD. However, our results show that PS is highly prone to manipulation for various combinations of agents and objects. This manipulation and the possible gain from manipulation become more severe particularly when agents are risk averse, and designers need to take this into consideration. On the other hand, while RSD does not theoretically guarantee *sd*-efficiency, our results show that it tends to do quite well – sometimes even outperforming PS in terms of social welfare. RSD also has the added advan-

tage of being *sd*-strategyproof and thus is not prone to the manipulation problems of PS.

Although computing RSD probabilities (fractional assignments) is #P-hard [6, 45], RSD is easy to implement in practice. However, the welfare cost of adopting manipulable mechanisms such as PS raises concern and has real consequences [16, 42]. Even though computing optimal manipulation strategies is computationally hard for the PS mechanism, evidentially individuals can easily figure out how to manipulate such mechanisms using simple greedy heuristics [16, 39]. Our investigations show that in many instances RSD performs as desirably as PS in terms of social welfare. Conversely, PS assignments are highly susceptible to manipulation especially when agents are risk averse.

These findings suggest that in multiagent settings where mechanism designers are unsure of sincere reporting of their preferences or when agents are mostly risk averse, the use of RSD is more desirable to ensure truthful reporting while providing reasonable social welfare. However, PS is still a desirable allocation mechanism for its fairness and efficiency properties, particularly in settings where agents are sincere.

References

1. Abdulkadiroğlu, A., Pathak, P.A., Roth, A.E.: Strategy-proofness versus efficiency in matching with indifference: redesigning the new york city high school match. Tech. rep., National Bureau of Economic Research (2009)
2. Abdulkadiroğlu, A., Sönmez, T.: Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* **66**(3), 689–701 (1998)
3. Abdulkadiroğlu, A., Sönmez, T.: House allocation with existing tenants. *Journal of Economic Theory* **88**(2), 233–260 (1999)
4. Arrow, K.J.: *Essays in the theory of risk-bearing* (1974)
5. Ashlagi, I., Fischer, F., Kash, I.A., Procaccia, A.D.: Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior* (2013)
6. Aziz, H., Brandt, F., Brill, M.: The computational complexity of random serial dictatorship. *Economics Letters* **121**(3), 341–345 (2013)
7. Aziz, H., Chen, J., Filos-Ratsikas, A., Mackenzie, S., Mattei, N.: Egalitarianism of random assignment mechanisms. arXiv preprint arXiv:1507.06827 (2015)
8. Aziz, H., Gaspers, S., Mackenzie, S., Mattei, N., Narodytska, N., Walsh, T.: Equilibria under the probabilistic serial rule. In: *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI 2015*, pp. 1105–1112. AAAI Press (2015). URL <http://dl.acm.org/citation.cfm?id=2832249.2832402>
9. Aziz, H., Gaspers, S., Mackenzie, S., Mattei, N., Narodytska, N., Walsh, T.: Manipulating the probabilistic serial rule. In: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, pp. 1451–1459. International Foundation for Autonomous Agents and Multiagent Systems (2015)
10. Aziz, H., Gaspers, S., Mackenzie, S., Walsh, T.: Fair assignment of indivisible objects under ordinal preferences. *Artificial Intelligence* **227**, 71–92 (2015)
11. Berg, S.: Paradox of voting under an urn model: the effect of homogeneity. *Public Choice* **47**(2), 377–387 (1985)
12. Bhargat, A., Chakrabarty, D., Khanna, S.: Social welfare in one-sided matching markets without money. In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pp. 87–98. Springer (2011)
13. Bogomolnaia, A., Heo, E.J.: Probabilistic assignment of objects: Characterizing the serial rule. *Journal of Economic Theory* **147**(5), 2072–2082 (2012)
14. Bogomolnaia, A., Moulin, H.: A new solution to the random assignment problem. *Journal of Economic Theory* **100**(2), 295–328 (2001)

15. Bouveret, S., Lang, J.: Manipulating picking sequences. In: In Proceedings of the 21st European Conference on Artificial Intelligence (ECAI14), pp. 141–146. IOS Press, Prague, Czech Republic (2014). URL <http://recherche.noiraudes.net/resources/papers/ECAI14.pdf>
16. Budish, E., Cantillon, E.: The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *The American economic review* **102**(5), 2237–71 (2012)
17. Che, Y.K., Kojima, F.: Asymptotic equivalence of probabilistic serial and random priority mechanisms. *Econometrica* **78**(5), 1625–1672 (2010)
18. Domshlak, C., Hüllermeier, E., Kaci, S., Prade, H.: Preferences in AI: An overview. *Artificial Intelligence* **175**(7), 1037–1052 (2011)
19. Ekici, Ö., Kesten, O.: An equilibrium analysis of the probabilistic serial mechanism. *International Journal of Game Theory* pp. 1–20 (2015). DOI 10.1007/s00182-015-0475-9. URL <http://dx.doi.org/10.1007/s00182-015-0475-9>
20. Filos-Ratsikas, A., Frederiksen, S.K.S., Zhang, J.: Social welfare in one-sided matchings: Random priority and beyond. In: *Algorithmic Game Theory*, pp. 1–12. Springer (2014)
21. Fishburn, P.C.: Lexicographic orders, utilities and decision rules: A survey. *Management Science* pp. 1442–1471 (1974)
22. Hadar, J., Russell, W.R.: Rules for ordering uncertain prospects. *The American Economic Review* pp. 25–34 (1969)
23. Hosseini, H., Larson, K.: Strategyproof quota mechanisms for multiple assignment problems. arXiv preprint arXiv:1507.07064 (2015)
24. Hugh-Jones, D., Kurino, M., Vanberg, C.: An experimental study on the incentives of the probabilistic serial mechanism. Tech. rep., Discussion Paper, Social Science Research Center Berlin (WZB), Research Area Markets and Politics, Research Unit Market Behavior (2013)
25. Kendall, M.G.: A new measure of rank correlation. *Biometrika* **30**(1/2), 81–93 (1938)
26. Kendall, M.G.: Rank correlation methods (1948)
27. Kojima, F.: Random assignment of multiple indivisible objects. *Mathematical Social Sciences* **57**(1), 134–142 (2009)
28. Kojima, F., Manea, M.: Incentives in the probabilistic serial mechanism. *Journal of Economic Theory* **145**(1), 106–123 (2010)
29. Liu, Q., Pycia, M.: Ordinal efficiency, fairness, and incentives in large markets. Unpublished mimeo (2013)
30. Lu, T., Boutilier, C.: Learning mallows models with pairwise preferences. In: *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 145–152 (2011)
31. Lu, T., Boutilier, C.: Robust approximation and incremental elicitation in voting protocols. In: *In Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI 2011*, vol. 22, p. 287 (2011)
32. Mallows, C.L.: Non-null ranking models. *Biometrika* **44**(1/2), 114–130 (1957)
33. Manea, M.: Asymptotic ordinal inefficiency of random serial dictatorship. *Theoretical Economics* **4**(2), 165–197 (2009)
34. Manlove, D.: *Algorithmics of matching under preferences*. World Scientific Publishing (2013)
35. Marden, J.I.: *Analyzing and modeling rank data*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability (1996)
36. Mattei, N., Walsh, T.: Preflib: A library of preference data [HTTP://PREFLIB.ORG](http://preflib.org). In: *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)*, Lecture Notes in Artificial Intelligence. Springer (2013)
37. McLennan, A.: Ordinal efficiency and the polyhedral separating hyperplane theorem. *Journal of Economic Theory* **105**(2), 435–449 (2002)
38. Mennle, T., Seuken, S.: Hybrid mechanisms: Trading off strategyproofness and efficiency in one-sided matching. Tech. rep. (2014). Working Paper
39. Mennle, T., Weiss, M., Philipp, B., Seuken, S.: The power of local manipulation strategies in assignment mechanisms. In: *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI’15*, pp. 82–89. AAAI Press (2015). URL <http://dl.acm.org/citation.cfm?id=2832249.2832261>
40. Pápai, S.: Strategyproof multiple assignment using quotas. *Review of Economic Design* **5**(1), 91–105 (2000)
41. Pathak, P.A.: Lotteries in student assignment. Unpublished mimeo, Harvard University (2006)

42. Pathak, P.A., Sethuraman, J.: Lotteries in student assignment: An equivalence result. *Theoretical Economics* **6**(1), 1–17 (2011)
43. Roth, A.E., Sönmez, T., Ünver, M.U.: Kidney exchange. *The Quarterly Journal of Economics* **119**(2), 457–488 (2004)
44. Saban, D., Sethuraman, J.: A note on object allocation under lexicographic preferences. *Journal of Mathematical Economics* (2013)
45. Saban, D., Sethuraman, J.: The complexity of computing the random priority allocation matrix. *Mathematics of Operations Research* **40**(4), 1005–1014 (2015)
46. Schulman, L.J., Vazirani, V.V.: Allocation of divisible goods under lexicographic preferences. *arXiv preprint arXiv:1206.4366* (2012)
47. Sönmez, T., Ünver, M.U.: Course bidding at business schools. *International Economic Review* **51**(1), 99–123 (2010)
48. Svensson, L.G.: Strategy-proof allocation of indivisible goods. *Social Choice and Welfare* **16**(4), 557–567 (1999)
49. Von Neumann, J.: A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions to the Theory of Games* **2**, 5–12 (1953)

Appendices

A Numerical Results

The following table shows the results of comparing RSD and PS under ordinal preferences for various combinations of agents and objects. Note that in most instances, RSD and PS do not induce the same random allocation.

n	m	Equal	Dominance		RSD	PS manipulability		
			SD	LD	weakEnvy	weak	SD	LD
2	2	100%	0%	0%	0%	0%	0%	0%
2	3	27%	18%	29%	23%	31%	31%	31%
2	4	10%	36%	60%	20%	53%	53%	53%
2	5	3%	39%	78%	16%	78%	78%	78%
2	6	1%	45%	90%	13%	87%	87%	87%
2	7	0%	46%	95%	12%	95%	95%	95%
2	8	0%	45%	96%	11%	97%	97%	97%
2	9	0%	47%	96%	11%	100%	100%	100%
2	10	0%	48%	99%	9%	99%	99%	99%
3	2	100%	0%	0%	0%	0%	0%	0%
3	3	67%	0%	0%	11%	24%	0%	0%
3	4	3%	5%	40%	47%	77%	5%	5%
3	5	0%	4%	75%	46%	96%	26%	27%
3	6	0%	6%	84%	42%	95%	53%	54%
3	7	0%	5%	90%	41%	100%	68%	69%
3	8	0%	5%	93%	39%	100%	80%	83%
3	9	0%	9%	96%	35%	100%	90%	92%
3	10	0%	7%	95%	34%	100%	94%	94%
4	2	62%	38%	38%	0%	0%	0%	0%
4	3	33%	34%	46%	21%	42%	0%	0%
4	4	21%	3%	8%	27%	72%	0%	0%
4	5	0%	0%	48%	61%	96%	1%	1%
4	6	0%	0%	76%	62%	98%	17%	18%
4	7	0%	0%	84%	62%	100%	33%	35%
4	8	0%	1%	93%	61%	99%	52%	54%
4	9	0%	1%	94%	60%	100%	65%	69%
4	10	0%	2%	95%	56%	100%	79%	85%
5	2	39%	61%	61%	0%	0%	0%	0%
5	3	8%	34%	83%	27%	66%	0%	0%
5	4	3%	19%	53%	42%	94%	0%	0%
5	5	6%	1%	7%	42%	90%	0%	0%
5	6	0%	0%	58%	69%	100%	0%	0%
5	7	0%	0%	84%	71%	100%	4%	4%
5	8	0%	0%	91%	71%	100%	18%	18%
5	9	0%	0%	94%	71%	100%	32%	36%
5	10	0%	0%	97%	70%	100%	49%	55%

Continued on the next page

n	m	Equal	Dominance		RSD	PS manipulability		
			SD	LD	weakEnvy	weak	SD	LD
6	2	21%	79%	79%	0%	0%	0%	0%
6	3	2%	71%	96%	31%	59%	0%	0%
6	4	0%	22%	88%	52%	90%	0%	0%
6	5	0%	9%	46%	59%	98%	0%	0%
6	6	3%	1%	7%	54%	96%	0%	0%
6	7	0%	0%	62%	74%	100%	0%	0%
6	8	0%	0%	89%	74%	100%	1%	1%
6	9	0%	0%	95%	75%	100%	8%	9%
6	10	0%	0%	97%	75%	100%	23%	25%
7	2	12%	88%	88%	0%	0%	0%	0%
7	3	1%	64%	99%	33%	83%	0%	0%
7	4	0%	26%	97%	57%	99%	0%	0%
7	5	0%	8%	87%	66%	100%	0%	0%
7	6	0%	2%	41%	69%	100%	0%	0%
7	7	1%	1%	6%	61%	99%	0%	0%
7	8	0%	0%	71%	79%	100%	0%	0%
7	9	0%	0%	93%	79%	100%	0%	0%
7	10	0%	0%	96%	78%	100%	5%	6%
8	2	8%	92%	92%	0%	0%	0%	0%
8	3	0%	63%	100%	34%	76%	0%	0%
8	4	0%	33%	99%	60%	95%	0%	0%
8	5	0%	10%	97%	70%	100%	0%	0%
8	6	0%	4%	83%	74%	100%	0%	0%
8	7	0%	1%	29%	74%	100%	0%	0%
8	8	0%	0%	5%	69%	99%	0%	0%
8	9	0%	0%	70%	81%	100%	0%	0%
8	10	0%	0%	93%	82%	100%	0%	0%
9	2	3%	97%	97%	0%	0%	0%	0%
9	3	0%	76%	100%	35%	70%	0%	0%
9	4	0%	33%	100%	62%	100%	0%	0%
9	5	0%	19%	99%	72%	100%	0%	0%
9	6	0%	6%	98%	76%	100%	0%	0%
9	7	0%	2%	78%	78%	100%	0%	0%
9	8	0%	0%	26%	78%	100%	0%	0%
9	9	0%	0%	4%	71%	100%	0%	0%
9	10	0%	0%	69%	84%	100%	0%	0%
10	2	2%	99%	99%	0%	0%	0%	0%
10	3	0%	70%	100%	37%	79%	0%	0%
10	4	0%	46%	100%	63%	98%	0%	0%
10	5	0%	17%	100%	73%	97%	0%	0%
10	6	0%	10%	99%	77%	100%	0%	0%
10	7	0%	2%	95%	79%	100%	0%	0%
10	8	0%	1%	77%	80%	100%	0%	0%
10	9	0%	0%	21%	79%	100%	0%	0%
Continued on the next page								

			Dominance		RSD	PS manipulability		
n	m	Equal	SD	LD	weakEnvy	weak	SD	LD
10	10	0%	0%	4%	73%	100%	0%	0%

Table 6: Experimental results over the space of preference profiles. SD (respectively LD) refers to the fraction of profiles where PS Stochastically (Lexicographically) Dominates RSD, and weakEnvy shows the average fraction of agents that are weakly envious under RSD. The last three columns show the fraction of profiles that PS is weakly manipulable, *sd*-manipulable (SD), and *ld*-manipulable (LD).